

Sparse Phase Retrieval: Convex Algorithms and Limitations

Kishore Jaganathan

Samet Oymak

Babak Hassibi

Department of Electrical Engineering
California Institute of Technology
Pasadena, CA 91125

Abstract—We consider the problem of recovering signals from their power spectral density. This is a classical problem referred to in literature as the phase retrieval problem, and is of paramount importance in many fields of applied sciences. In general, additional prior information about the signal is required to guarantee unique recovery as the mapping from signals to power spectral density is not one-to-one. In this paper, we assume that the underlying signals are sparse.

Recently, semidefinite programming (SDP) based approaches were explored by various researchers. Simulations of these algorithms strongly suggest that signals upto $o(\sqrt{n})$ sparsity can be recovered by this technique. In this work, we develop a tractable algorithm based on reweighted l_1 -minimization that recovers a sparse signal from its power spectral density for significantly higher sparsities, which is unprecedented. We also discuss the limitations of the existing SDP algorithms, and provide a combinatorial algorithm which requires significantly fewer measurements to guarantee recovery.

Index Terms—Phase Retrieval, Semidefinite Programming (SDP), Reweighted l_1 -minimization

I. INTRODUCTION

In many practical measurement systems, the power spectral density of the signal, i.e. the magnitude square of the Fourier transform, is the measurable quantity. Phase information of the Fourier transform is completely lost, because of which signal recovery is difficult. This problem occurs in many areas of engineering and applied physics, including X-ray crystallography [3], astronomical imaging [4], microscopy, optics [5], blind channel estimation and so on.

Recovering a signal from its Fourier transform magnitude, or equivalently its autocorrelation, is known as phase retrieval. The mapping from signals to their Fourier transform magnitude is not one-to-one, and hence unique recovery is not possible in general. Additional measurements or prior information about the signal is required in order to uniquely recover the underlying signal. Constraints on the signal's values like non-negativity, bounds on the signal's support (locations where the value is non-zero), and more recently sparsity [9], are commonly used as prior information.

This work was supported in part by the National Science Foundation under grants CCF-0729203, CNS-0932428 and CCF-1018927, by the Office of Naval Research under the MURI grant N00014-08-1-0747, and by Caltech's Lee Center for Advanced Networking.

Considerable amount of research has been done over the last few decades ([1], [2]) and a wide range of heuristics have been proposed (see [6]), a comprehensive survey of which can be found in [7]. [8] provides a theoretical framework to understand the heuristics, which are in essence an alternating projection between a convex set and a non-convex set. Such methods often converge to a local minimum, hence drastically reducing the chances of successful signal recovery.

Recently, the phase retrieval problem was recast as a semi-definite programming problem (see [10], [11] and [12]). In [10], additional measurements with different illuminations, which is possible in an optical setup, are used to make unique recovery feasible. In [11], [12] and [13], the underlying signals are assumed to be sparse, which is a reasonable assumption in applications like X-ray crystallography, microscopy and astronomical imaging.

Numerical simulations of the existing techniques based on SDP strongly suggest that signals upto $o(\sqrt{n})$ sparsity can be recovered with an arbitrarily high probability. This behavior for the phase retrieval problem was rigorously explained in [17], [18]. [19], [20] consider the "generalized" phase retrieval problem and observe a similar behavior. In this work, we develop an algorithm based on the idea of reweighted l_1 minimization to solve the phase retrieval problem for significantly higher sparsities. We also provide certain theoretical guarantees and discuss the limitations of the SDP based techniques, and develop a combinatorial technique which requires far fewer measurements to guarantee recovery.

The remainder of the paper is organized as follows. In Section 2, we formulate the phase retrieval problem and recast it as a SDP problem. In Section 3, we discuss the limitations of the existing SDP-based techniques and develop an algorithm based on reweighted l_1 minimization in Section 4. In Section 5, we develop a measurement system using a combinatorial approach which requires far fewer measurements to guarantee recovery. Section 6 presents the results of the numerical simulations and concludes the paper.

II. PROBLEM FORMULATION

Suppose $\mathbf{x} = (x_0, x_1, \dots, x_{n-1})$ is a real-valued discrete-time signal of length n and sparsity k , where sparsity is defined as the number of non-zero entries. Let $\mathbf{y} = (y_0, y_1, \dots, y_{n-1})$

be its Fourier transform, i.e.,

$$\mathbf{y} = \mathbf{F}\mathbf{x} \quad (1)$$

where \mathbf{F} is the $n \times n$ DFT matrix. The phase retrieval problem can be mathematically stated as

$$\begin{aligned} &\text{find} && \mathbf{x} \\ &\text{subject to} && |\mathbf{y}| = |\mathbf{F}\mathbf{x}| \end{aligned} \quad (2)$$

Since magnitude square of Fourier transform and autocorrelation are Fourier pairs, the phase retrieval problem can be reformulated as recovering signals from their autocorrelation, i.e.,

$$\begin{aligned} &\text{find} && \mathbf{x} \\ &\text{subject to} && a_i = \sum_j x_j x_{j+i} \quad 0 \leq i \leq n-1 \end{aligned} \quad (3)$$

where $\mathbf{a} = (a_0, a_1, \dots, a_{n-1})$ is the autocorrelation of \mathbf{x} .

Observe that the operations of time-shift, flipping and global sign-change do not affect the autocorrelation, because of which there is a trivial ambiguity. The signals resulting from these operations are considered equivalent, and in all the applications it is considered good enough if any equivalent signal is recovered.

The problem (2) is hard to solve due to non-convex constraints. We can relax the constraints into a set of convex constraints by embedding the problem in a higher dimensional space, a technique popularly known as lifting. Note that (2) contains n constraints of the form $|y_i| = |\mathbf{f}_i^T \mathbf{x}|$, where \mathbf{f}_i is the i^{th} column of \mathbf{F} . Squaring both sides, the constraints can be rewritten as $|y_i|^2 = |\mathbf{x}^T \mathbf{f}_i^* \mathbf{f}_i^T \mathbf{x}|$. Suppose we define $\mathbf{X} = \mathbf{x}\mathbf{x}^T$, the problem can be recast in terms of \mathbf{X} as

$$\begin{aligned} &\text{find} && \mathbf{X} \\ &\text{subject to} && |\mathbf{y}_i|^2 = \text{trace}(\mathbf{M}_i \mathbf{X}) \quad 0 \leq i \leq n-1 \\ &&& \text{rank}(\mathbf{X}) = 1 \quad \& \quad \mathbf{X} \succeq 0 \end{aligned} \quad (4)$$

where $\mathbf{M}_i = \mathbf{f}_i^* \mathbf{f}_i^T$.

In terms of the autocorrelation \mathbf{a} , the lifted problem can be formulated as

$$\begin{aligned} &\text{find} && \mathbf{X} \\ &\text{subject to} && \sum_{j=1}^n \mathbf{X}_{j,j+i} = a_i \quad 0 \leq i \leq n-1 \\ &&& \text{rank}(\mathbf{X}) = 1 \quad \& \quad \mathbf{X} \succeq 0 \end{aligned} \quad (5)$$

III. BACKGROUND

The program (4) can be reformulated as a rank minimization problem as follows

$$\begin{aligned} &\text{minimize} && \text{rank}(\mathbf{X}) \\ &\text{subject to} && |\mathbf{y}_i|^2 = \text{trace}(\mathbf{M}_i \mathbf{X}) \quad 0 \leq i \leq n-1 \\ &&& \mathbf{X} \succeq 0 \end{aligned} \quad (6)$$

This is a non-convex problem as rank is a non-convex function. It has been shown in [14] that trace minimization is the

tightest convex relaxation of rank minimization for positive semidefinite matrices. This relaxation is not useful in the phase retrieval setup as $\text{trace}(\mathbf{X})$ corresponds to the energy of the signal \mathbf{x} , which is fixed by the magnitude of the Fourier transform. [15] proposes log-determinant function as a surrogate for rank in such problems, i.e.,

$$\text{minimize} \quad \log \det (\mathbf{X} + \epsilon \mathbf{I}) \quad (7)$$

$$\text{subject to} \quad |\mathbf{y}_i|^2 = \text{trace}(\mathbf{M}_i \mathbf{X}) \quad 0 \leq i \leq n-1 \quad (8)$$

$$\mathbf{X} \succeq 0 \quad (9)$$

This heuristic tries to minimize a concave function in a convex domain, which can be done using gradient descent approach. This method was explored for the phase retrieval setup in [10], [11]. Simulations suggest that the algorithm converges to a rank 1 solution with high probability if the underlying signal is $o(\sqrt{n})$ sparse.

In [12], we explored a two-stage recovery process to provably solve (2). In the first stage, we use information about the support of the autocorrelation to recover the support of the signal (see [17]). In the second stage, we solve the SDP with known support [18]. It was empirically observed and theoretically shown that signals were recovered with arbitrarily high probability if the sparsity was $o(\sqrt{n})$. However, if the support information was available by other means, it was observed that the program recovered signals up to roughly $o(n)$ sparsity.

IV. RECOVERY ALGORITHM

In this section, we develop an iterative algorithm based on reweighted l_1 -minimization to solve the phase retrieval problem outside the $o(\sqrt{n})$ sparsity regime.

A two-stage approach like [12] wouldn't work as the support of the autocorrelation becomes full if the signal has sparsity greater than $O(\sqrt{n \log(n)})$. Trace minimization in the phase retrieval setup has two issues: trivial ambiguities have same objective function, trace is fixed because of which we will be solving a feasibility problem only. Weighted l_1 minimization (10) intuitively overcomes these issues and promotes sparse solutions.

$$\begin{aligned} &\text{minimize} && \text{trace}(\mathbf{V}|\mathbf{X}|) \\ &\text{subject to} && |\mathbf{y}_i|^2 = \text{trace}(\mathbf{M}_i \mathbf{X}) \quad 0 \leq i \leq n-1 \\ &&& \mathbf{X} \succeq 0 \end{aligned} \quad (10)$$

where \mathbf{V} is a weight matrix, which can be designed to promote the necessary structure in the solution.

Simulations suggest that (10) has rank 1 solutions with high probability if the sparsity is $o(\sqrt{n})$ and \mathbf{V} is chosen randomly, but fails outside the $o(\sqrt{n})$ region. However, the largest eigenvalue turns out to be considerably stronger than the other eigenvalues, and the eigenvector corresponding to it happens to contain a lot of information about the support of the signal even significantly outside the $o(\sqrt{n})$ region, which is not very surprising.

This strongly suggests the possibility of an iterative algorithm, which at every iteration also knows "a lot" about where the signal's non-zero entries can be. Algorithm 1 uses this information by doing a re-weighted minimization at every iteration. The weights corresponding to prospective support locations are set to zero to encourage the signal to choose those locations in the next iteration, and the weights outside this region is chosen to be positive.

Algorithm 1 Phase Retrieval Algorithm

Input: The magnitude of the Fourier transform $|\mathbf{y}|$, maximum number of iterations

Output: The underlying signal \mathbf{x}

1. Initialize \mathbf{V} by choosing its entries from $[0, 1]$ uniformly at random
2. Solve the optimization problem

$$\begin{aligned} & \text{minimize} && \text{trace}(\mathbf{V}|\mathbf{X}|) \\ & \text{subject to} && |\mathbf{y}_i|^2 = \text{trace}(\mathbf{M}_i \mathbf{X}) \quad 0 \leq i \leq n-1 \\ & && \mathbf{X} \succeq 0 \end{aligned} \quad (11)$$

3. If $\text{rank}(\mathbf{X}) = 1$, return \mathbf{X} , else calculate $\mathbf{X}_1 = \mathbf{x}_1 \mathbf{x}_1^T$, where \mathbf{X}_1 is the best rank-1 approximation of \mathbf{X}
 4. Update \mathbf{V} as follows: $\mathbf{V}_{ij} = 0$ if $|\mathbf{x}_i|$ and $|\mathbf{x}_j|$ are greater than a certain threshold, choose the remaining entries from $[0, 1]$ uniformly at random
 5. Iterate until convergence or maximum number of iterations
 6. Calculate \mathbf{X}_1 and return \mathbf{x}_1
-

V. PHASE RETRIEVAL AS A COMPRESSION PROBLEM

Theorem V.1 was proved for the phase retrieval problem in [17], [18].

Theorem V.1. *Signals can be recovered from their power spectral densities up to time-shift, reversal and global sign with probability $1 - \delta$ for any $\delta > 0$ if*

- 1) $k = o(\sqrt{n})$
- 2) k entries are chosen uniformly at random
- 3) $n > n(\delta)$

We can instead fix the sparsity k and consider the phase retrieval problem with partial power spectral density information (for example [22]). In particular, one might have access to only certain frequencies. To solve this problem, we will use some classical results in the compressed sensing literature.

Theorem V.2 (Candes, Romberg, Tao, [23]). *Let \mathbf{F} be the DFT matrix. Consider the random matrix \mathbf{A} obtained by choosing m rows of \mathbf{F} uniformly at random. If \mathbf{x} is a k sparse vector, \mathbf{x} can be recovered from observations $\mathbf{A}\mathbf{x}$ with arbitrarily high probability if $m \geq O(k \log n)$ via the following ℓ_1 minimization:*

$$\begin{aligned} & \min_{\hat{\mathbf{x}}} && \|\hat{\mathbf{x}}\|_1 \\ & \text{subject to} && \mathbf{A}\hat{\mathbf{x}} = \mathbf{A}\mathbf{x} \end{aligned} \quad (12)$$

The following theorem gives a useful result for the recovery of a sparse signal from its partial power spectral density.

Theorem V.3. *Let \mathbf{x} be a k sparse vector whose support is chosen uniformly at random. Suppose we observe its power spectral density at m distinct frequencies chosen uniformly at random. If $m \geq O(k^2 \log n)$, \mathbf{x} can be recovered from its partial power spectral density with arbitrarily high probability.*

Proof: If \mathbf{x} is a k -sparse signal, its autocorrelation can have at most k^2 non-zero entries. Since power spectral density is the Fourier transform of the autocorrelation, using Theorem V.2, whenever $m \geq O(k^2 \log n)$, the autocorrelation of \mathbf{x} can be recovered from partial power spectral density observations via ℓ_1 minimization with arbitrarily high probability. ■

While Theorem V.3, gives a tractable algorithm for the recovery of \mathbf{a} from partial power spectral density observations $\mathbf{s} = \mathbf{A}\mathbf{a}$, one can consider a more sophisticated approach. The program (12) tries to solve for a k^2 -sparse solution without using the extra information that the resulting signal should be a valid autocorrelation. We propose the following method which might be of interest for future directions.

$$\begin{aligned} & \min_{\mathbf{a}} && \|\mathbf{a}\|_1 \end{aligned} \quad (13)$$

$$\text{subject to} \quad \mathbf{A}\mathbf{a} = \mathbf{s}$$

$$\begin{aligned} & \sum_{j=1}^n \mathbf{X}_{j,j+1} = a_i \quad 0 \leq i \leq n-1 \\ & \text{rank}(\mathbf{X}) = 1 \quad \& \quad \mathbf{X} \succeq 0 \end{aligned} \quad (14)$$

Observe that the only nonconvex constraint in (13) is (14). We believe that solving (13) can substantially increase the performance and instead of $O(k^2 \log n)$ measurements, just $O(k \log n)$ measurements might suffice which is similar to that of typical compressed recovery of a k sparse vector. However, further investigation is required to relax (14) in a useful manner as opposed to dropping it.

Overall, Theorem V.3 is subject to a $o(\sqrt{n})$ -sparsity bottleneck since the full power spectral density corresponds to n measurements. While we do not provide theoretical guarantees for Algorithm 1, when full power spectral density is available, our algorithm seemingly beats the $o(\sqrt{n})$ -sparsity bottleneck. In section VI, we see that the recoverable sparsity is much higher than $o(\sqrt{n})$.

A. Two-stage recovery

In [12], we try to solve the sparse phase retrieval problem via a two-stage approach where the first step involves finding the support of the signal from the support of the autocorrelation. We will now argue that such an approach is inherently subject to the $o(\sqrt{n})$ bottleneck, as there is no way of finding the support of the signal when its sparsity is greater than $O(\sqrt{n})$.

Lemma V.1. *Suppose \mathbf{x} is a k -sparse signal whose support is chosen uniformly at random, and whose nonzero entries are continuous i.i.d. random variables. Then, there exists a*

constant c such that whenever $k \geq c\sqrt{n\log(n)}$, support of the autocorrelation is full with arbitrarily high probability.

Proof: Without loss of generality, we can assume that each location belongs to the support of the signal with probability $\frac{k}{n}$ independently as the same proof will apply for k -sparse signals with standard modifications.

For a particular distance d , if no two non-zero entries in the signal are separated by a distance d , we can say that d doesn't belong to the support of the autocorrelation. This probability can be bounded by $(1 - k^2/n^2)^{n/2}$ which is upper bounded by $e^{-k^2/2n}$. Union bound tells us that the probability of the support of the autocorrelation not being full is less than $ne^{-k^2/2n}$, which goes to zero if $k \geq c\sqrt{n\log(n)}$ for sufficiently large n . ■

B. Relation to Gaussian Phase Retrieval

Our results on partial power spectral density can be related to the "generalized" phase retrieval problem, where the observations are of the form $|\mathbf{g}_i^T \mathbf{x}|$ for i.i.d. complex standard normal vectors $\{\mathbf{g}_i\}_{i=1}^m$. While this problem is structurally similar to phase retrieval, it is considerably simpler as there are no trivial ambiguities like time-shift and flipping.

Assuming \mathbf{x} is a sparse vector, [19] and [20] analyze the following semidefinite program for the tractable recovery of \mathbf{x} up to a global phase ambiguity.

$$\begin{aligned} \min_{\hat{\mathbf{x}}} \quad & \|\hat{\mathbf{x}}\|_* + \lambda \|\hat{\mathbf{x}}\|_1 \\ \text{subject to} \quad & \langle \mathbf{g}_i \mathbf{g}_i^T, \hat{\mathbf{x}} \rangle = \langle \mathbf{g}_i \mathbf{g}_i^T, \mathbf{x} \mathbf{x}^T \rangle \quad 1 \leq i \leq m \end{aligned} \quad (15)$$

Naturally, one would wish that the unique minimizer of (15) to be $\mathbf{x} \mathbf{x}^T$ to be able to recover \mathbf{x} up to phase ambiguity. (15) tries to capture both sparsity and low rankness of the underlying matrix $\mathbf{x} \mathbf{x}^T$. Interestingly, both [19] and [20] suggest that as long as $m \leq O(\min\{k^2, n\})$, recovery using (15) is not possible with very high probability for any choice of regularizer λ . To summarize, even if $k^2 \ll n$, one would still need $\Omega(k^2)$ measurements for recovery, which indicates an $o(\sqrt{n})$ bottleneck for generic Gaussian measurements too.

On the other hand, [19], [20] and [21] show that, if one is able to search for a low rank and sparse matrix, then, $\mathbf{x} \mathbf{x}^T$ can be recovered with only $k \log \frac{n}{k}$ measurements which illustrates a significant performance gap between tractable recovery and intractable combinatorial search. Interestingly, [12] illustrates a similar gap for the phase retrieval problem with Fourier measurements.

Overall, we have seen that for the two important class of phaseless measurements (Fourier and Gaussian), the current theoretical guarantees for the existing algorithms are subject to a strong $o(\sqrt{n})$ -sparsity bottleneck. A natural question is whether it is possible at all to do sparse phase retrieval in a tractable way with small number of measurements.

C. Combinatorial Approach

To address this question, we consider the problem of recovering \mathbf{x} from phaseless measurements by making use of a

specific choice of measurements. In particular, we show that one can tractably recover a k -sparse signal from phaseless measurements by using only $O(k \log n)$ measurements with very high probability. This shows that phase retrieval with optimal number of measurements (up to a $\log n$ factor) is in fact possible.

Theorem V.4. Suppose \mathbf{x} is an arbitrary k -sparse vector where $k \geq 2$. Let $\{\mathbf{z}_i\}_{i=1}^m$ be i.i.d. vectors with i.i.d. entries distributed as:

$$\begin{cases} 0 & \text{with probability } 1 - 1/k \\ \mathcal{N}(0, 1) & \text{with probability } 1/k \end{cases} \quad (16)$$

Let $\{\mathbf{a}_i, \mathbf{b}_i\}_{i=1}^m$ be i.i.d. vectors with i.i.d. entries distributed as $\exp(i\theta)$ where θ is uniformly distributed in $[0, 2\pi)$. Denote the function $\mathbf{R}^n \times \mathbf{R}^n \rightarrow \mathbf{R}^n$ that returns entrywise products of two vectors by \cdot . Assume, we observe the measurements:

$$|\langle \mathbf{a}_i \cdot \mathbf{z}_i, \mathbf{x} \rangle|^2, |\langle \mathbf{b}_i \cdot \mathbf{z}_i, \mathbf{x} \rangle|^2 \quad (17)$$

for $1 \leq i \leq m$.

Then, $\mathbf{x} \mathbf{x}^*$ can be recovered with high probability in $O(mn)$ time, whenever $m \geq ck \log n$ for some constant $c > 0$.

Remark: In general, we don't need the knowledge of sparsity k for the design of the measurement operator. This can be handled by introducing an extra factor of $\log n$ in the required number of measurements, i.e. $k \log^2(n)$ measurements. The reason is we can take $ck \log n$ measurements designed for each of the sparsity levels of $k_i = 2^i$ for $1 \leq i \leq \lceil \log n \rceil$, i.e. use k_i instead of k in (16). Then, k will lie between k_i and k_{i+1} for some i and the same proof argument will work.

The proof of Theorem V.4 can be found in the appendix. While our proposed recovery algorithm requires a perfectly sparse signal, we emphasize that, our intention with Theorem V.4 is demonstrating the possibility of tractable recovery rather than coming up with a robust algorithm for artificially designed measurements (16).

VI. NUMERICAL SIMULATIONS

In order to demonstrate the performance of Algorithm 1, numerical simulations were performed for different values of signal length n and sparsities k .

For a given n and sparsity k , the k support locations were chosen from the n locations uniformly at random. The signal values at the support locations were chosen from $[0, 1]$ uniformly at random. The Fourier transform magnitudes of the signal were computed and provided as input to Algorithm 1. If the output signal matched the input signal, it was counted as a success, else it was counted as a failure. Two other semidefinite program based algorithms (CandesPR: [10], HassibiPR: [12]) and the traditional alternating projection algorithm (GS: [6]) were used for comparison purposes.

The results of the numerical simulations are shown in Figure 1. The probability of successful recovery is plotted against various sparsities for $n = 32$ and $n = 64$. 100 simulations

were performed for each sparsity to get the average success probability, which was calculated as the percentage of successful recovery. The existing SDP based algorithms start fading at around $o(\sqrt{n})$ sparsity, the figure clearly demonstrates the superiority of using reweighted minimization outside the $o(\sqrt{n})$ sparsity region.




Fig. 1. Success rate of recovery for $n = 32$ and various sparsities




Fig. 2. Success rate of recovery for $n = 64$ and various sparsities

REFERENCES

- [1] A. L. Patterson: A direct method for the determination of the components of interatomic distances in crystals, *Zeitschr. Krist.* 90 (1935) 517-542
- [2] A.L. Patterson: Ambiguities in the X-ray analysis of crystal structures, *Phys. Review* 65 (1944) 195-201.
- [3] R. P. Millane, "Phase retrieval in crystallography and optics," *J. Opt. Soc. Am. A* 7, 394-411 (1990)
- [4] J.C. Dainty and J.R. Fienup, "Phase Retrieval and Image Reconstruction for Astronomy," Chapter 7 in H. Stark, ed., *Image Recovery: Theory and Application* pp. 231-275.
- [5] A. Walther, "The question of phase retrieval in optics," *Opt. Acta* 10, 4149 (1963).
- [6] R. W. Gerchberg and W. O. Saxton. "A practical algorithm for the determination of the phase from image and diffraction plane pictures". *Optik* 35, 237 (1972).

- [7] J. R. Fienup, "Phase retrieval algorithms: a comparison". *Appl. Opt.* 21, 2758-2769 (1982).
- [8] H.H. Bauschke, P.L. Combettes and D.R. Luke, "Phase retrieval, error reduction algorithm, and Fienup variants: a view from convex optimization" *J. Opt. Soc. Am.A* (2002).
- [9] Y.M. Lu and M. Vetterli. "Sparse spectral factorization: Unicity and reconstruction algorithms". *ICASSP* 2011.
- [10] E. J. Candes, Y. Eldar, T. Strohmer and V. Voroninski, "Phase retrieval via matrix completion", *arXiv:1109.0573 [cs.IT]*, 2011.
- [11] Y. Shechtman, Y. C. Eldar, A. Szameit, and M. Segev, "Sparsity Based Sub-Wavelength Imaging with Partially Incoherent Light Via Quadratic Compressed Sensing", *Optics Express*, vol. 19, Issue 16, pp. 14807-14822, Aug. 2011.
- [12] K. Jaganathan, S. Oymak, and B. Hassibi, "Recovery of Sparse 1-D Signals from the Magnitudes of their Fourier Transform". *Information Theory Proceedings (ISIT)*, 2012 IEEE International Symposium on , vol., no., pp.1473,1477, 1-6 July 2012.
- [13] Y. Shechtman, A. Beck, Y. C. Eldar, "GESPAR: Efficient Phase Retrieval of Sparse Signals", *arXiv:1301.1018*.
- [14] B. Recht, M. Fazel, and P. Parrilo. "Guaranteed Minimum-Rank Solutions of Linear Matrix Equations via Nuclear Norm Minimization". *SIAM Review*, Vol 52, no 3.
- [15] M. Fazel, H. Hindi, and S. Boyd. "Log-det Heuristic for Matrix Rank Minimization with Applications to Hankel and Euclidean Distance Matrices". *Proc. American Control Conference*, 2003.
- [16] J. Chen, A. Ghosh, J. Magutt, and M. Chiang. "QAVA: quota aware video adaptation" In *Proceedings of the 8th international conference on Emerging networking experiments and technologies (CoNEXT '12)*.
- [17] K. Jaganathan and B. Hassibi, "Reconstruction of Integers from Pairwise Distances", *arXiv:1212.2386 [cs.DM]*.
- [18] K. Jaganathan, S. Oymak and B. Hassibi, "On Robust Phase Retrieval for Sparse Signals", *Allerton Conference on Communication, Control and Computing*, 2012.
- [19] S. Oymak, A. Jalali, M. Fazel, Y. C. Eldar, and B. Hassibi, "Simultaneously Structured Models with Application to Sparse and Low-rank Matrices", *arXiv:1212.3753*.
- [20] X. Li, V. Voroninski, "Sparse Signal Recovery from Quadratic Measurements via Convex Programming", *arXiv:1209.4785v1*.
- [21] Y. C. Eldar and S. Mendelson, "Phase Retrieval: Stability and Recovery Guarantees", *arXiv:1211.0872*.
- [22] M. Moravec, J. Romberg and R. Baraniuk, "Compressed Sensing Phase Retrieval", *SPIE, Wavelets XII*, vol. 6701, 2007.
- [23] E. J. Candes, J. Romberg and T. Tao. "Stable signal recovery from incomplete and inaccurate measurements." *Comm. Pure Appl. Math.*, 59 1207-1223.
- [24] P. Erdős, A. Rényi (1960). "On the evolution of random graphs". *Publ. Math. Inst. Hung. Acad. Sci., Ser. A* 5, 17-61 (1960).
- [25] M. L. Fredman, D. E. Willard, "Trans-dichotomous algorithms for minimum spanning trees and shortest paths", *Journal of Computer and System Sciences* 48 (3): 533-551.

Proof of Theorem V.4

Proof: The proof provides a tractable recovery algorithm and consists of two main steps.

Support recovery: Denote the support of \mathbf{x} by $S \subseteq [n]$ and the support of \mathbf{z}_i by S_i and let $\mathbf{c}_i = \mathbf{z}_i \cdot \mathbf{a}_i$. Now, consider the inner product $\mathbf{c}_i^* \mathbf{x}$. Since nonzero entries of \mathbf{c}_i have continuous distribution if $S \cap S_i \neq \emptyset$, $|\mathbf{c}_i^* \mathbf{x}| \neq 0$ almost surely. Hence whenever a measurement $|\mathbf{c}_i^* \mathbf{x}|^2 = 0$ we can deduce that $S \cap S_i = \emptyset$. Let:

$$I = \{1 \leq i \leq m \mid S_i \cap S = \emptyset\} \quad (18)$$

Clearly, $\{S_i\}$'s are i.i.d. supports and for each i we have:

$$\mathbb{P}(S \cap S_i = \emptyset) = (1 - 1/k)^k \geq \frac{1}{4} \quad (19)$$

Following (19), by the law of large numbers, w.h.p. $|I| \geq m/8$. Conditioned on $|I| \geq m/8$, the probability that $j \in \bar{S}$ is not contained in $\bigcup_{i \in I} S_i$ is at most $(1 - 1/k)^{|I|} \leq (1 - 1/k)^{m/8}$.

Assuming $m \geq 8k \log n$ and using a union bound:

$$\mathbb{P}(\bar{S} \not\subseteq \bigcup_{i \in I} S_i) \leq \exp(-\frac{m}{8k}) \leq n^{-1} \quad (20)$$

which will approach 0. Hence, the exact support of the signal can be found by simply taking the union of sets S_i satisfying $\mathbf{c}_i^* \mathbf{x} = 0$ and then complementing it ($O(mn)$ time). Next, with the knowledge of support, we proceed with the recovery of \mathbf{x} up to an overall phase ambiguity.

Signal recovery: Recovery of the signal given its support will be performed in two steps. We first show that magnitudes of nonzero entries of \mathbf{x} can be found.

Recovering magnitudes: Assume $S_i \cap S$ is a singleton $j \in S$. Since we already have the knowledge of S from previous part, we can immediately deduce: $|x_j|^2 = \frac{|\langle \mathbf{c}_i, \mathbf{x} \rangle|^2}{|\mathbf{c}_{i,j}|^2}$. Then, we simply need to ensure that w.h.p., for all $j \in S$ there exists $1 \leq i \leq m$ satisfying $S \cap S_i = j$. Probability of this not happening for a fixed $j \in S$ is:

$$(1 - \frac{1}{k}(1 - \frac{1}{k})^{k-1})^m \leq \exp(-\frac{m}{4k}). \quad (21)$$

After union bounding over all $j \in S$, whenever $m \geq 8k \log n$, all $|x_j|^2$ can be found with probability at least $1 - k \exp(-\frac{m}{4k}) \geq 1 - n^{-1}$.

Recovering the relative phases: Next, we consider the measurements satisfying $|S_i \cap S| = 2$, $1 \leq i \leq m$. For fixed i we have:

$$\mathbb{P}(|S_i \cap S| = 2) = \binom{k}{2} \frac{1}{k^2} (1 - \frac{1}{k})^{k-2} \geq \frac{1}{8}. \quad (22)$$

Overall, w.h.p. there are $m/10$ measurements satisfying $|S_i \cap S| = 2$. Let:

$$I = \{1 \leq i \leq m \mid |S_i \cap S| = 2\} \quad (23)$$

Next, form the k vertex graph obtained by connecting the nodes j, l whenever $\{j, l\} = S_i \cap S$ for some $i \in I$ ($O(mn)$

time). Observe that, each $i \in I$ picks an edge in this graph uniformly at random. Overall, we perform at least $\frac{m}{10}$ picks with replacement. This graph is connected with high probability whenever the chance of an edge being picked is more than $\frac{2 \log k}{k}$, [24]. This happens when $m \geq ck \log k$ for $c \geq 10$ since:

$$\mathbb{P}(\text{edge is picked}) \geq 1 - (1 - \frac{1}{\binom{k}{2}})^{m/10} \quad (24)$$

$$\geq 1 - \exp(-\frac{m}{5k(k-1)}) \quad (25)$$

$$\geq 1 - \exp(-\frac{c \log k}{5k}) \approx \frac{c \log k}{5k} \quad (26)$$

Now, w.h.p. the graph is connected. Find a spanning tree T in this graph, which can be done in $O(k^2)$ time, [25]. Set the phase of the initial node to 0. Then, the rest of the phases of the nodes are uniquely determined as follows.

Recovering relative phases: To begin, consider an edge of T between nodes j, l where $\{j, l\} = S_i \cap S$ for some i . Fix $x_j = |x_j|$ i.e. 0 phase. From measurements:

$$|\langle \mathbf{a}_i \cdot \mathbf{z}_i, \mathbf{x} \rangle|^2, |\langle \mathbf{b}_i \cdot \mathbf{z}_i, \mathbf{x} \rangle|^2 \quad (27)$$

and with the knowledge of $|x_i|^2, |x_j|^2$, we can find: $\Re(z_j z_l a_j^* a_l x_j x_l^*)$ and $\Re(z_j z_l b_j^* b_l x_j x_l^*)$ where $\Re(\cdot)$ returns the real part of a number. This information is equivalent to: $\Re(ax), \Re(bx)$ where $a = a_j^* a_l$, $b = b_j^* b_l$ have uniformly random phases and $x = x_j x_l^*$. We will argue that the phase of x is uniquely determined when $\Re(ax), \Re(bx), a, b, |x|$ are known. Assume $a = \exp(i\theta_1)$, $b = \exp(i\theta_2)$ and $x = |x| \exp(i\theta)$. Then,

$$\Re(ax) = |x| \cos(\theta + \theta_1) \quad (28)$$

$$\Re(bx) = |x| \cos(\theta + \theta_2) \quad (29)$$

gives us two linearly independent equations (almost surely) and two unknowns ($\sin \theta, \cos \theta$). Overall, this would yield θ . Applying this over all edges of the tree recursively, we can find the exact phase differences between all neighboring nodes and since the graph is connected we can find the phase of any $j \in S$ by adding up the phase differences over all edges on the path between the initial node and j . This can be done for all nodes in $O(k^2)$ time via DFS. Overall, the original signal \mathbf{x} can be found up to an overall phase ambiguity in $O(mn)$ time. ■